

Feature Article
Supercomputing Affects Everyone
November 2006

For many people, the word “supercomputer” often brings to mind a room-size machine with complex controls and a dazzling display of lights. It seems totally alien to most people, because it has no part in their daily lives. This could not be further from the truth.

Today, the majority of the population uses computers in some manner. There are computer chips embedded in televisions, cars, games, watches, DVDs, etc. Many of these embedded computer chips are just older models of the chips in desktop or laptop PCs. Most people don't realize that the **technologies in the average home computer were originally pioneered in the supercomputing field**. Let us look at some of the technologies that have historically come from the supercomputing realm into everyday use.

Graphics

In the early 1980s, home computers were not common. In those days, big names in the business like Atari, Sinclair, Osborne, and Kaypro, were known only to microcomputer specialists. The screens on these computers displayed blocky letters in various shades of green. Even the first computers to bear the names Macintosh and IBM PC had low resolution, single-color displays. At that time, being able to display top of the line 3D color graphics meant purchasing machines with names like IRIS, Indigo, or Onyx from a company named Silicon Graphics Inc. (SGI to its friends). These machines cost 10s to 100s of thousands of dollars or more. SGI made graphics workstations and graphics servers with massive video boards that were covered in chips. These chips rendered graphics using SGI's proprietary protocol named GL for “Graphics Language”. Today, computer gaming enthusiasts can purchase reasonably priced computers with 3D graphics rendering capability that is integrated into a single chip on the video card. These 3D video cards are rendering graphics using OpenGL, a standard derived from the original GL language. **High-end graphics is just one technology that has made its way from supercomputers into the desktop computer.**

Pipelining and Vector Processing

Another change has been in the way CPUs (central processing units) handle multiple pieces of data. Early CPUs were single-cycle processors. This means that in a single clock tick the CPU would fetch the instruction, decode the instruction, execute it, and write back the results. A better way to do this is to design the CPU to use pipelining. This means that one part of the CPU is fetching instruction 1, while another part is decoding instruction 2, and another

part is executing instruction 3, etc. This pipelining idea is very much like an automobile assembly line, where many cars are on the line at different stages of assembly, instead of building one car at a time. Another variation on this theme is vector processing, in which an array of numbers is held in memory and a mathematical operation is performed on all the numbers at once. Vector processing requires having multiple arithmetic logic units (ALU). **These chips, designed to handle multiple pieces of data at once, were originally pioneered for use in supercomputers. Today the CPUs and video rendering chips in desktop computers utilize similar designs.**

Parallel Processing

The idea of working with multiple pieces of data at once can be taken one step further. Along with a single chip handling multiple pieces of data, why not have multiple CPUs working on the same problem at the same time? Supercomputers have been designed to work with multiple CPUs for several decades. This in turn has led to an enormous effort to design programming tools that make it possible to write software that uses multiple CPUs at once. In the high performance computing industry, this is called parallel processing. In the past 10 years, it became possible to get multiple CPUs in mid range computers, designed to be servers or workstations. In order to make multiple CPUs affordable for home and desktop machines, AMD and Intel have begun putting two CPUs on a single computer chip. These are called dual core chips. For some months now it has been possible to purchase affordable desktop computers with dual core CPUs such as the AMD Athlon64 X2 chips, and more recently Intel Core 2 Duo chips. At present, software manufacturers are scrambling to redesign their software to utilize parallel processing. Thus, the **programming tools developed for supercomputing applications are now being applied to developing commodity software for home and office computers.**

FPGAs

This discussion can be taken a bit further. To get an idea of what types of things might be coming to home computers in a few years, the cutting edge of supercomputing can be examined. One technology on the cutting edge of supercomputing is the FPGA chip (Field Programmable Gate Array). The chips in desktop computers have a fixed set of circuits for executing instructions. An FPGA chip can be reconfigured to have different circuits for every piece of software that runs on it. Often FPGA programmers will set up the FPGA chip to have hundreds of circuits to do a key task. This makes a computer with an FPGA chip run with the speed of a computer with many CPUs. Imagine how fast your computer would run if you had a hundred special World of Warcraft chips, or a hundred PowerPoint chips. FPGAs are just gaining momentum in the supercomputing realm, so it will probably be some years before they are seen in home computers.

The Alabama Supercomputer Center will be installing FPGA chips in the Cray XD1 supercomputer in late 2006. Software for writing VHDL code to run on FPGA chips will be installed at the same time. VHDL stands for Very High Speed Integrated Circuit Hardware Description Language. It is the low level language used for programming FPGAs. Unlike machine code, which describes a set of instructions, VHDL describes electronic circuits. Users will also have access to Impulse C, a C-like language compiler for FPGAs. **The Cray XD1 is one of the few computer platforms designed to utilize FPGAs in a production-computing environment.**

Most software packages have not yet been adapted to run on FPGAs. However, researchers need access to this type of technology if they are to be able to write software that does utilize these chips effectively. **The Alabama Supercomputer Authority has chosen to invest in this technology as part of their role as a high performance computing technology leader in the state of Alabama.** Students that learn to utilize such cutting edge resources are developing skills that will help them negotiate career challenges in the future.

Contact

If you have questions about using the supercomputers at the Alabama Supercomputer Center, please contact David Young at dyoung@asc.edu or 256-971-7434.